

# Towards a practical parallelisation of the simplex method

J. A. J. Hall

15<sup>th</sup> February 2005

## Abstract

The simplex method is frequently the most efficient method of solving general large sparse linear programming (LP) problems. Despite this, there has been no parallelisation of the simplex method that offers significantly improved performance over a good serial implementation of the revised simplex method. This paper reviews previous attempts to parallelising the simplex method and finds them to be, at worst, hopelessly inefficient and, at best, of insufficient value for general problems to be adopted practically. However, encouraging ideas are identified, pointing the way to worthwhile future developments in this area.

**Keywords:** linear programming, simplex method, sparse, parallel computing

**MSC classification:** 90C05

## 1 Introduction

Linear programming (LP) is a widely applicable technique both in its own right and as a sub-problem in the solution of other optimization problems. The revised simplex method and the barrier method are the two efficient methods for solving general large sparse LP problems. In a context where families of related LP problems have to be solved, such as in integer programming and decomposition methods, and for certain classes of single LP problems, the revised simplex method is usually the more efficient method.

The application of parallel and vector processing to the simplex method for linear programming has been considered since the early 1970's. However, only since the beginning of the 1980's have attempts been made to develop implementations, with the period from the late 1980's until the mid to late 1990's seeing the greatest activity. Although there have been a few experiments using vector processing, the vast majority of implementations have made use of distributed memory multiprocessors and ethernet-connected clusters. For such architectures, the high ratio of communication to computation, in particular when using efficient techniques for the revised simplex method, has been a major factor in the lack of success. In shared memory multiprocessors, coordination of parallel processes may be achieved very much more efficiently by memory instructions. It is, therefore, surprising that since the late 1990's, when the availability of shared memory multiprocessors has become increasingly widespread, there has been little work on parallelising the simplex method. The time is, perhaps, ripe for further work in this area. This paper reviews the simplex method, its variants and their implementation in detail, with an eye to developing efficient parallel implementations.

Section 2 introduces the simplex method and discusses issues of implementation and behaviour that influence its parallelisation. Terms and concepts in parallel computing that are used in this paper are introduced briefly in Section 3. A review of parallel simplex and simplex-like methods for solving LP problems is given in Section 4. As an outcome of this review, several promising approaches for future parallel implementations are identified in Section 5.

## 2 The simplex method

The simplex method and its computational requirements are most conveniently discussed in the context of LP problems in standard form

$$\begin{aligned} & \text{minimize} && \mathbf{c}^T \mathbf{x} \\ & \text{subject to} && A\mathbf{x} = \mathbf{b} \\ & && \mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{1}$$

where  $\mathbf{x} \in \mathbb{R}^n$  and  $\mathbf{b} \in \mathbb{R}^m$ . Although LP problems may be formulated with general bounds on variables and inequality constraints, such problems may be transformed into the form (1). Conversely, the solution techniques discussed in this paper for problems of the form (1) may be modified to solve more general LP problems without any significant computational overhead. The matrix  $A$  in (1) may well contain columns of the identity corresponding

to *logical* (slack) variables introduced to transform inequality constraints into equations. The remaining columns of  $A$  correspond to *structural* (original) variables.

In the simplex method, the indices of variables are partitioned into sets  $\mathcal{B}$  corresponding to  $m$  basic variables  $\mathbf{x}_B$ , and  $\mathcal{N}$  corresponding to  $n - m$  nonbasic variables  $\mathbf{x}_N$ , such that the basis matrix  $B$  formed from the columns of  $A$  corresponding to  $\mathcal{B}$  is nonsingular. The set  $\mathcal{B}$  itself is conventionally referred to as the basis. The columns of  $A$  corresponding to  $\mathcal{N}$  form the matrix  $N$  and the components of  $\mathbf{c}$  corresponding to  $\mathcal{B}$  and  $\mathcal{N}$  are referred to as, respectively, the basic costs  $\mathbf{c}_B$  and non-basic costs  $\mathbf{c}_N$ .

When the nonbasic variables are set to zero the values  $\hat{\mathbf{b}} = B^{-1}\mathbf{b}$  of the basic variables, if non-negative, correspond to a vertex of the feasible region. The expression  $\mathbf{x}_B + B^{-1}N\mathbf{x}_N = \hat{\mathbf{b}}$  derived from the equations in (1) allows the basic variables to be eliminated from the objective which becomes  $(\mathbf{c}_N^T - \mathbf{c}_B^T B^{-1}N)\mathbf{x}_N + \mathbf{c}_B^T \hat{\mathbf{b}}$ . If none of the components of the vector of *reduced costs*  $\hat{\mathbf{c}}_N = \mathbf{c}_N^T - \mathbf{c}_B^T B^{-1}N$  is positive then the current basis is optimal.

In each iteration of the simplex method, if the current basis is not optimal a nonbasic variable  $x_q$  with negative reduced cost is chosen to enter the basis. Increasing this variable from zero whilst maintaining the equations in (1) corresponds to moving along an edge of the feasible region such that the objective function decreases. The direction of this edge is given by the column  $\hat{\mathbf{a}}_q$  of  $\hat{N} = B^{-1}N$  corresponding to  $x_q$ . By considering the ratios of the components of  $\hat{\mathbf{b}}$  to the corresponding components of  $\hat{\mathbf{a}}_q$  (when positive), the simplex method finds the first basic variable to be reduced to zero as  $x_q$  is increased and, hence, the step to the next vertex of the feasible region along this edge. Repeating this process determines a path along edges of the feasible region which terminates at a vertex for which there is either no edge along which the objective decreases (so the vertex is optimal) or the objective may be reduced without bound along an edge (so the LP problem is unbounded).

There are many strategies for choosing the nonbasic variable  $x_q$  to enter the basis. The original rule of choosing the variable with the most negative reduced cost is commonly referred to as the Dantzig criterion. However, if the components of  $\hat{\mathbf{a}}_j$  are large relative to  $\hat{c}_j$  it is likely that only a small increase in  $x_j$  will be possible before one of the basic variables is reduced to zero. Alternative pricing strategies weight the reduced cost by dividing it by (a measure of) the length of  $\hat{\mathbf{a}}_j$ . The exact steepest edge strategy, described by Goldfarb and Reid [32], maintains weights  $s_j = 1 + \|\hat{\mathbf{a}}_j\|^2$  corresponding to the step length for a unit change in  $x_j$ . The *Devex* strategy due to Harris [39] maintains approximate edge weights.

A popular technique for choosing the variable to leave the basis is the EXPAND procedure of Gill *et al.* [30]. By expanding the constraint bounds by a small amount, this strategy often enables the leaving variable to be chosen from a set of possibilities on grounds of numerical stability. The technique is of particular relevance in the commonly-encountered situation of degeneracy, although it is not guaranteed to prevent cycling, as shown by Hall and McKinnon [37].

For a given basis, if any basic variables are negative they are given costs of  $-1$ , with other variables being given costs of zero. Phase I iterations are then performed in order to force all variables to be non-negative, if possible. When this is achieved, the simplex method reverts to the original problem costs and determines an optimal vertex or unboundedness using the process, referred to as phase II, introduced above.

The two main variants of the simplex method correspond to different means of calculating the data required to determine the step to the new vertex. The first variant is the standard simplex method in which the reduced costs and the directions of all edges at the current vertex are maintained in a rectangular tableau. In the revised simplex method, the reduced costs and the direction of the chosen edge are determined by solving systems involving the basis matrix  $B$ .

## 2.1 The standard simplex method

In the standard simplex method the matrix  $\hat{N}$ , the (reduced) right-hand-side vector  $\hat{\mathbf{b}}$ , the reduced costs  $\hat{\mathbf{c}}_N$  and current value of the objective  $\hat{f} = \mathbf{c}_B^T \hat{\mathbf{b}}$  are maintained in a tableau (2).

	$\mathcal{N}$	RHS	
$\mathcal{B}$	$\hat{N}$	$\hat{\mathbf{b}}$	(2)
	$\hat{\mathbf{c}}_N^T$	$-\hat{f}$	

Each iteration of the standard simplex method requires a Gauss-Jordan elimination operation to be applied to the columns of the tableau so that the updated tableau corresponds to the new basis.

The simplex method is commonly started from a basis for which  $B = I$  so the matrix in the standard simplex tableau is  $N$ . As such the tableau is sparse. It is widely assumed that the extent of fill-in caused by the elimination operations is such that it is not worth exploiting sparsity. As a consequence, the standard simplex method is generally implemented using a dense data structure.

## 2.2 The revised simplex method

The computational components of the revised simplex method are illustrated in Figure 1. At the beginning of an iteration, it is assumed that the vector of reduced costs  $\hat{\mathbf{c}}_N^T$  and the vector  $\hat{\mathbf{b}}$  of current values of the basic variables are known and that a representation of  $B^{-1}$  is available. The first operation is CHUZC which scans the weighted reduced costs to determine a good candidate  $q$  to enter the basis. The *pivotal column*  $\hat{\mathbf{a}}_q$  is formed using the representation of  $B^{-1}$  in an operation referred to as FTRAN.

The CHUZR operation determines the variable to leave the basis, with  $p$  being used to denote the index of the row in which the leaving variable occurred, referred to as the *pivotal row*. The index of the leaving variable itself is denoted by  $p'$ . Once the indices  $q$  and  $p'$  have been interchanged between the sets  $\mathcal{B}$  and  $\mathcal{N}$ , a *basis change* is said to have occurred. The RHS vector  $\hat{\mathbf{b}}$  is then updated to correspond to the increase  $\alpha = \hat{b}_p / \hat{a}_{pq}$  in  $x_q$ .

```

CHUZC: Scan  $\hat{\mathbf{c}}_N$  for a good candidate  $q$  to enter the basis.
FTRAN: Form the pivotal column  $\hat{\mathbf{a}}_q = B^{-1}\mathbf{a}_q$ , where  $\mathbf{a}_q$  is column  $q$  of  $A$ .
CHUZR: Scan the ratios  $\hat{b}_i / \hat{a}_{iq}$  for the row  $p$  of a good candidate to leave the
        basis. Let  $\alpha = \hat{b}_p / \hat{a}_{pq}$ .
        Update  $\hat{\mathbf{b}} := \hat{\mathbf{b}} - \alpha \hat{\mathbf{a}}_q$ .
BTRAN: Form  $\boldsymbol{\pi}_p^T = \mathbf{e}_p^T B^{-1}$ .
PRICE: Form the pivotal row  $\hat{\mathbf{a}}_p^T = \boldsymbol{\pi}_p^T N$ .
        Update reduced costs  $\hat{\mathbf{c}}_N^T := \hat{\mathbf{c}}_N^T - \hat{c}_q \hat{\mathbf{a}}_p^T$ .
If {growth in representation of  $B$ } then
    INVERT: Form a new representation of  $B^{-1}$ .
else
    UPDATE: Update the representation of  $B^{-1}$  corresponding to the basis
            change.
end if

```

Figure 1: Operations in an iteration of the revised simplex method

Before the next iteration can be performed it is necessary obtain the reduced costs and a representation of the new matrix  $B^{-1}$ . Although the reduced costs may be computed directly using the following operations,

$$\begin{aligned} \boldsymbol{\pi}_B^T &= \mathbf{c}_B^T B^{-1} \\ \hat{\mathbf{c}}_N^T &= \mathbf{c}_N^T - \boldsymbol{\pi}_B^T N, \end{aligned} \tag{3}$$

it is more efficient computationally to update them by computing the pivotal row  $\hat{\mathbf{a}}_p^T = \mathbf{e}_p^T B^{-1} N$  of the standard simplex tableau. This is obtained in two

steps. First the vector  $\pi_p^T = e_p^T B^{-1}$  is formed using the representation of  $B^{-1}$  in an operation known as BTRAN, and then the vector  $\hat{a}_p^T = \pi_p^T N$  of values in the pivotal row is formed. This sparse matrix-vector product with  $N$  is referred to as PRICE. Once the reduced costs have been updated, the UPDATE operation modifies the representation of  $B^{-1}$  according to the basis change. Note that, periodically, it will generally be either more efficient, or necessary for numerical stability, to find a new representation of  $B^{-1}$  using the INVERT operation.

During phase I, Gill *et al.* [30] describe how the EXPAND procedure for CHUZR is modified, although there is no significant change in its computational cost. The only significant computational requirement in phase I which is not indicated in Figure 1 occurs when one or more of the variables which remain basic become feasible so their cost coefficients increase by one to zero. In order to update the reduced costs, it is necessary to compute the corresponding linear combination of tableau rows. This composite row is formed by the following BTRAN and PRICE operations

$$\begin{aligned}\pi_\delta^T &= \delta^T B^{-1} \\ \hat{a}_\delta^T &= \pi_\delta^T N,\end{aligned}$$

where the nonzeros in  $\delta$  are equal to one for each variable which remains basic and becomes feasible.

When using the Devex strategy [39], if the pivotal row is computed to update the reduced costs then it may be used to update the Devex weights at no significant computational cost. To update the exact steepest edge weights requires, in addition to the pivotal row, a further BTRAN operation to form  $\hat{a}_q^T B^{-1}$  and a further PRICE operation to form the product of this vector with  $N$ . It is also computationally expensive to initialise steepest edge weights if the initial basis matrix is not the identity. As a consequence of these overheads, and since the Devex strategy performs well in terms of reducing the number of iterations required to solve LP problems, Devex is commonly the default in efficient sequential implementations of the revised simplex method.

### 2.3 The revised simplex method with multiple pricing

When in-core memory was severely restricted, a popular variant of the revised simplex method was to incorporate minor iterations of the standard simplex method, restricted to a small subset of the variables. This is described by Orchard-Hays [59] and is referred to as *multiple pricing*. The major computational steps of this variant, together with those required when using

Devex or steepest edge weights, are illustrated in Figure 2. The features of these variants, in particular those incorporating Devex or steepest edge, make exploiting parallelism particularly attractive.

The CHUZC operation scans the (weighted) reduced costs to determine a set  $\mathcal{Q}$  of good candidates to enter the basis. The inner loop then applies the standard simplex method to the LP problem corresponding to the candidates in  $\mathcal{Q}$  so requires the corresponding columns  $\hat{\mathbf{a}}_j = B^{-1}\mathbf{a}_j$ ,  $j \in \mathcal{Q}$ , of the standard simplex tableau. These columns are formed as a multiple FTRAN. The matrix formed by the columns  $\hat{\mathbf{a}}_j$ ,  $j \in \mathcal{Q}$  and the corresponding vector of reduced costs for the candidates  $j \in \mathcal{Q}$  are conveniently denoted by  $\hat{\mathbf{a}}_{\mathcal{Q}}$  and  $\hat{\mathbf{c}}_{\mathcal{Q}}$ .

In each minor iteration, CHUZC\_MI scans the (weighted) reduced costs of the candidates in  $\mathcal{Q}$  and selects one,  $q$  say, to enter the basis. Once a basis change has been determined by CHUZR and the RHS has been updated, the standard simplex tableau corresponding to the new basis is obtained by updating the previous tableau in an operation known as UPDATE\_MI. The matrix  $\hat{\mathbf{a}}_{\mathcal{Q}}$  and reduced costs  $\hat{\mathbf{c}}_{\mathcal{Q}}$  are updated by the standard Gauss-Jordan elimination step. Any Devex weights are updated using row  $p$  of the updated tableau. Any steepest edge weights may be computed directly using the known tableau columns. Minor iterations are terminated when either  $\mathcal{Q} = \emptyset$  or there are no negative reduced costs for  $q \in \mathcal{Q}$ .

An original advantage of multiple pricing was that the reduced costs were only computed (from scratch) every major iteration using 3. As a result, the frequency with which BTRAN, and particularly PRICE, were performed was reduced significantly.

Leaving aside any advantages due to efficient use of memory, the value of the revised simplex method with multiple pricing depends on the extent to which the variables that were good candidates to enter the basis when  $\mathcal{Q}$  was formed remain good candidates, or even remain attractive, during the course of minor iterations. This property is commonly referred to as *candidate persistence* and its effect depends upon both the problem being solved and the number of candidates chosen when forming  $\mathcal{Q}$ . Since the entering variable in the second and subsequent minor iteration is not necessarily the best candidate according to the original selection criterion, it is possible that the number of basis changes required to solve the LP problem will increase, perhaps significantly. If any of the original candidates becomes unattractive and minor iterations terminate with  $\mathcal{Q} \neq \emptyset$ , the FTRAN and update operations required to form and maintain the tableau column for such a candidate are wasted.

CHUZYC: Scan  $\hat{\mathbf{c}}_N$  for a set  $\mathcal{Q}$  of good candidates to enter the basis.  
 FTRAN: Form  $\hat{\mathbf{a}}_j = B^{-1}\mathbf{a}_j, \forall j \in \mathcal{Q}$ , where  $\mathbf{a}_j$  is column  $j$  of  $A$ .  
 Loop {minor iterations}  
   CHUZYC\_MI: Scan  $\hat{\mathbf{c}}_Q$  for a good candidate  $q$  to enter the basis.  
   CHUZR: Scan the ratios  $\hat{b}_i/\hat{a}_{iq}$  for the row  $p$  of a good candidate to leave the basis, where  $\hat{\mathbf{b}} = B^{-1}\mathbf{b}$ . Let  $\alpha = \hat{b}_p/\hat{a}_{pq}$ .  
   UPDATE\_MI: Update  $\mathcal{Q} := \mathcal{Q} \setminus \{q\}$ .  
               Update  $\hat{\mathbf{b}} := \hat{\mathbf{b}} - \alpha\hat{\mathbf{a}}_q$ .  
               Update the columns  $\hat{\mathbf{a}}_Q$  and reduced costs  $\hat{\mathbf{c}}_Q$ .  
   If {Devex} then  
     Update the Devex weights for the candidates in  $\mathcal{Q}$ .  
   else if {steepest edge} then  
     Update the steepest edge weights for the candidates in  $\mathcal{Q}$ .  
   end if  
 End loop {minor iterations}  
 For {each basis change} do  
   If {Devex} then  
     BTRAN: Form  $\boldsymbol{\pi}_p^T = \mathbf{e}_p^T B^{-1}$ .  
     PRICE: Form pivotal row  $\hat{\mathbf{a}}_p^T = \boldsymbol{\pi}_p^T N$ .  
            Update reduced costs  $\hat{\mathbf{c}}_N := \hat{\mathbf{c}}_N - \hat{c}_q\hat{\mathbf{a}}_p^T$  and Devex weights.  
   else if {steepest edge} then  
     BTRAN: Form  $\boldsymbol{\pi}_p^T = \mathbf{e}_p^T B^{-1}$ .  
     PRICE: Form pivotal row  $\hat{\mathbf{a}}_p^T = \boldsymbol{\pi}_p^T N$ .  
            Update reduced costs  $\hat{\mathbf{c}}_N := \hat{\mathbf{c}}_N - \hat{c}_q\hat{\mathbf{a}}_p^T$ .  
     BTRAN: Form  $\mathbf{w} = \hat{\mathbf{a}}_q^T B^{-1}$ .  
     PRICE: Form  $\mathbf{w}^T \mathbf{a}_j$  for nonzero components of pivotal row  $\hat{\mathbf{a}}_p^T$ .  
   end if  
   If {growth in factors} then  
     INVERT: Form a new representation of  $B^{-1}$ .  
   else  
     UPDATE: Update the representation of  $B^{-1}$  corresponding to the basis change.  
   end if  
 End do  
 If {Dantzig} then  
   BTRAN: Form  $\boldsymbol{\pi}_B^T = \mathbf{c}_B^T B^{-1}$ .  
   PRICE: Form  $\hat{\mathbf{c}}_N^T = \mathbf{c}_N^T - \boldsymbol{\pi}_B^T N$ .  
 end if

Figure 2: The revised simplex method with multiple pricing



## 2.4 The representation of $B^{-1}$

Forming and updating the explicit inverse matrix of  $B$  as a dense data structure is simple and, if reformed when appropriate, is numerically stable. The FTRAN and BTRAN operations thus reduce to dense matrix-vector products. However, for general large sparse LP problems, this approach is as inefficient as the standard simplex method. Efficient implementations of the revised simplex method use a data structure corresponding to a factored representation of  $B^{-1}$ . This representation is based on that obtained by the INVERT operation for some basis matrix  $B_0$ . There are various techniques for updating the factored representation of  $B^{-1}$ , the original and simplest of which being the product form update of Dantzig and Orchard-Hays [21].

### 2.4.1 Procedures for INVERT

Although there are many techniques for factorising a general unsymmetric sparse matrix, in the context of the revised simplex method, there is much commonality of approach. The general aim is to use Gaussian elimination to obtain a factored representation of  $B^{-1}$ , seeking a reconciliation of the partially conflicting goals of low computational cost, low fill-in and numerical stability. Using an active submatrix that is, initially,  $B$ , Gaussian elimination determines a sequence of  $m$  entries to be used as pivots. Each pivot is used to eliminate any other entries in the corresponding column. The active submatrix is then updated by removing the row and column corresponding to the pivot. No INVERT technique is optimal for all LP problems.

LP basis matrices frequently contain significant numbers of columns of the identity matrix corresponding to logical variables. Any efficient INVERT must exploit this at the very least. The number of nonzeros in such columns being unity, by pivoting on these nonzeros the entries in the corresponding rows can be eliminated without incurring any fill-in. Once such logical columns and their corresponding rows have been removed there may well be structural columns which then have only one nonzero (in rows of the active submatrix). Similarly there may be rows which contain only a single nonzero. Such rows and columns which have unit ‘count’ in rows and columns in the active submatrix are referred to as *singletons* and their nonzero may be used as a pivot without any fill-in. As singletons are identified, pivoted on and their corresponding rows and columns removed, further singletons may be created. This *triangularisation* phase ends when no further singletons can be identified. Any remaining rows and columns have count of at least two and form what is frequently referred to as the *bump*.

For certain LP problems, in particular network problems, the triangular-



<pre> if (<math>r_{p_k} \neq 0</math>) then   <math>r_{p_k} := \mu_k r_{p_k}</math>   <math>\mathbf{r} := \mathbf{r} - r_{p_k} \boldsymbol{\eta}_k</math> end if </pre>	$r_{p_k} := \mu_k (r_{p_k} - \mathbf{r}^T \boldsymbol{\eta}_k)$
(a) FTRAN	(b) BTRAN

Figure 3: Standard operations in FTRAN and BTRAN

### 2.4.2 Updating the factored representation of $B^{-1}$

If  $E_U$  is used to denote the transformation of  $B_0$  corresponding the basis changes following INVERT, such that  $B = B_0 E_U$ , it follows that  $B^{-1}$  may be expressed as

$$B^{-1} = E_U^{-1} B_0^{-1}.$$

The product form update leaves the factored form of  $B_0^{-1}$  unaltered and represents  $E_U^{-1}$  as a product of pairs of elementary matrices of the form (4). In the UPDATE operation, the modification of  $B^{-1}$  corresponding to the basis change is obtained directly from the pivotal column and is given by  $p_k = p$ ,  $\mu_k = 1/\hat{a}_{pq}$  and  $\boldsymbol{\eta}_k = \hat{\mathbf{a}}_q - \hat{a}_{pq} \mathbf{e}_p$ .

In solvers based on the product form, the representation of the UPDATE operations can be appended to the eta file following INVERT, resulting in a single homogeneous data structure. However, in this paper, the particular properties of the INVERT and UPDATE etas and the nature of the operations with them may be exploited, so FTRAN is viewed as the pair of operations

$$\tilde{\mathbf{a}}_q = B_0^{-1} \mathbf{a}_q \quad (\text{I-FTRAN})$$

followed by

$$\hat{\mathbf{a}}_q = E_U^{-1} \tilde{\mathbf{a}}_q. \quad (\text{U-FTRAN})$$

Conversely, BTRAN is viewed as

$$\tilde{\boldsymbol{\pi}}^T = \mathbf{r}^T E_U^{-1} \quad (\text{U-BTRAN})$$

followed by

$$\boldsymbol{\pi}^T = \tilde{\boldsymbol{\pi}}^T B_0^{-1}. \quad (\text{I-BTRAN})$$

Other procedures for UPDATE modify the factored representation of  $B_0^{-1}$  with the aim of reducing the size of the eta file and achieving greater numerical stability. Whilst they may be more efficient in serial for some problems, there is a significant value in leaving the factored representation of  $B_0^{-1}$  unaltered. Techniques for improving the efficiency of FTRAN and BTRAN, both

in serial and parallel, may require an analysis of the representation of  $B_0^{-1}$  following INVERT that is only a worthwhile investment if many FTRAN and BTRAN operations are performed with the original representation of  $B_0^{-1}$ .

## 2.5 Hyper-sparsity

For most sparse LP problems the pivotal column  $\hat{\mathbf{a}}_q$  has sufficient zero entries that it is efficient to exploit this when creating the corresponding product form eta in UPDATE. Similarly, it is often found that the vector  $\boldsymbol{\pi}_p^T = \mathbf{e}_p^T B^{-1}$  has relatively few nonzero entries. If PRICE is performed as a set of inner products  $\boldsymbol{\pi}_p^T \mathbf{a}_i$ ,  $i \in \mathcal{N}$ , then there will be many operations with zeros. In this case it is more efficient to evaluate  $\boldsymbol{\pi}_p^T N$  as a linear combination of the rows of  $N$  corresponding to the nonzeros in  $\boldsymbol{\pi}_p$ , even allowing for the overhead of updating a row-wise representation of  $N$  following each basis change.

For some important classes of LP problems, the vectors  $\hat{\mathbf{a}}_q$ ,  $\boldsymbol{\pi}_p$ , and pivotal tableau row  $\hat{\mathbf{a}}_p^T$  are sparse. This property, which Hall and McKinnon [36] refer to as *hyper-sparsity*, has also been identified by Bixby *et al.* [12, 13]. For such problems, Hall and McKinnon [38] have shown that significant performance improvement can be gained by exploiting hyper-sparsity. This is achieved by maintaining the indices of the nonzeros in these vectors, and exploiting this information, when they are both formed and used. For the basis matrices of LP problems that exhibit hyper-sparsity the dimension of the bump within INVERT is very small relative to that of the matrix.

For some LP problems  $B^{-1}$ , row  $p$  of which is  $\boldsymbol{\pi}_p^T$ , is typically sparse. This property has been exploited in some implementations of the revised simplex method based on an explicit inverse. However, it is general for the number of nonzeros in  $B^{-1}$  (even when freshly formed) to be vastly more than the number of nonzeros in an efficient factored representation. This property is illustrated, for example, by Shu [64].

## 2.6 PRICE when the column/row ratio is large

For LP problems with a large column/row ratio ( $n \gg m$ ), the cost of the revised simplex variants set out in Figures 1 and 2 is dominated by that of PRICE. However, for such problems, it is typical for a good candidate to enter the basis to be identified from only a small subset of the columns. The numerous variants of partial and multiple pricing [59] stem from this observation. In general, their use leads to a reduction in the solution time that is significant, and in some cases vast, despite the fact that they are largely incompatible with the use of Devex weights.

## 3 Parallel computing

In classifying the approaches to attempts to parallelise the simplex method, reference is made to terms and concepts in parallel computing. This section introduces the necessary terms. A full and more general introduction to parallel computing is given by Kuman *et al.* [47].

### 3.1 Parallel multiprocessors

Parallel multiprocessor architectures may be classified as being either *distributed memory*, when each processor has its own local memory, or *shared memory*, when all processors have access to a common memory. There are two principal parallel programming paradigms. If the work of a major algorithmic operation may be distributed across a set of processors then *data parallelism* may be exploited. Conversely, if it is possible to perform several major algorithmic operations simultaneously, then *task parallelism* may be exploited. In practice, it may be possible to exploit both data parallelism and task parallelism for a particular (set of) major algorithmic operation(s).

### 3.2 Speed-up, scalability and Amdahl's law

In general, success in parallelisation is measured in terms of *speed-up*, the time required to solve a problem with more than one parallel processor compared with the time required using a single processor. The traditional goal is to achieve a speed-up factor equal to the number of processors. Such a factor is referred to as *linear speed-up* and corresponds to a *parallel efficiency* of 100%. The increase in available cache, and frequently RAM, with the number of parallel processors occasionally leads to the phenomenon of *super-linear speed-up* being observed. Parallel schemes for which, at least in theory, performance improves linearly and without limit as the number of processors increases are said to be *scalable*. If parallelism is not exploited in all major algorithmic operations then the speed-up is limited, according to Amdahl's law [2], by the proportion of the execution time for the non-parallel operations.

### 3.3 Programming paradigms and their implementation

There are two fundamental means of implementing algorithms on parallel machines. On distributed memory machines the concept of communicating data between processors using instructions initiated by explicit calls to

*message-passing* subroutines is natural and immediate. On shared memory machines the conceptually natural technique is *data-parallel* programming in which operations are coded as if they were to be executed serially but translated into a parallel executable by use of an appropriate compiler. Many message-passing environments are also supported on shared memory machines and data-parallel programming techniques are also supported on distributed memory machines.

On distributed memory machines, the overhead of sending messages to communicate data between processors is determined by both *latency* and *bandwidth*. The former is a time overhead that is independent of the size of the message and the latter is the rate of communication. For generic message-passing environments, on a particular architecture the latency and bandwidth may be significantly higher than for an architecture-dependent environment that is usually supplied and tuned by the vendor. When an algorithm has a high ratio of communication to computation, growing communication overhead may outweigh any improvement in potential computational performance resulting from the use of an increased number of processors.

## 4 Parallel simplex and simplex-like methods

Despite its value in both the academic and commercial worlds, there has been no parallelisation of the simplex method that offers significantly improved performance over a good serial implementation of the revised simplex method for general large scale LP problems. The successes in terms of speed-up typically correspond to parallelisation of implementations that are partly or wholly inefficient in serial. In most cases these are inefficiencies relative to good serial revised simplex implementations *of the time*. However, some techniques that have been parallelised are only seen as being inefficient in the light of serial revised simplex techniques that were either little known at the time or developed subsequently. This is identified below to emphasise that, as a result of the huge increase in efficiency of serial revised simplex solvers during the period of research into parallel simplex, the task of developing a worthwhile parallel simplex-based solver has increased enormously.

Attention has mainly been directed at inefficient simplex and simplex-like techniques in which exploiting parallelism is relatively straightforward. There have been only a few discussions of schemes by which parallelism might be exploited by (variants of) the revised simplex method with a factored inverse. Fewer still of these schemes have been implemented and none has met with any real success.

## 4.1 Parallel simplex using dense matrix algebra

The dense standard simplex method and the revised simplex method with a dense explicit inverse have been implemented in parallel many times. The simple data structures involved and potential for linear speed-up makes them attractive exercises in parallel computing. However, for solving general large scale LP problems, the serial inefficiency of these implementations is such that only with a prohibitively large number of parallel processes could they conceivably compete with a good sparsity-exploiting serial implementation of the revised simplex method.

Some authors acknowledge the superiority of the revised simplex method when solving large sparse problems and justify the use of dense matrix algebra by identifying the relatively small number of applications where large dense LP problems occur. For general large sparse LP problems the computational cost of an iteration of the revised simplex method with a dense inverse is  $O(m^2)$ . This is comparable with the  $O(mn)$  cost of the standard simplex method, and so is of no greater practical interest.

Several authors have claimed the revised simplex method with a dense inverse to be valuable when it is applied to sparse problems with high column/row ratios ( $n \gg m$ ) since the sparsity and parallelism in PRICE are easy to exploit. When this is done, the cost per iteration is vastly less. Indeed, for such problems, a parallel implementation of the revised simplex method with a dense inverse could be competitive with a serial implementation of the revised simplex method with a sparse factored inverse if the latter did not make use of partial or multiple pricing.

Early work on parallelising the simplex method using dense matrix algebra was mainly restricted to discussion of data distribution and communication schemes, with implementations limited to small numbers of processes on distributed memory machines. A survey of this work is given by Luo *et al.* [51] and other early references to work in this area are due to Finkel [27], Boffey and Hay [16], Babayev and Mardanov [8] and Agrawal *et al.* [1]. A relatively early work which is representative of this era is due to Stunkel [65] who implemented both the dense standard simplex method and the revised simplex method with a dense inverse on a 16-processor Intel hypercube, achieving a speed-up of between 8 and 12 for small problems from the Netlib set [29]. Cvetanovic *et al.* [19] report a speed-up of 12 when solving two small problems using the standard simplex method, a result that is notable for being achieved on a 16-processor *shared* memory machine. Luo and Reijns [50] obtained speed-ups of more than 12 on 16 transputers when using the revised simplex method with a dense inverse to solve modest Netlib problems. However, they developed their parallel scheme on the assumption

that the computational cost of updating the inverse using the expression  $B^{-1} := E_k^{-1}B^{-1}$  was  $O(m^3)$ , rather than  $O(m^2)$  if the structure of  $E_k^{-1}$  (4) is properly exploited, an error which might well explain their high efficiency.

Eckstein *et al.* [24] parallelised the standard simplex method and the revised simplex method with a dense inverse on the massively parallel Connection Machine CM-2 and CM-5, incorporating the steepest-edge pricing strategy [32] in their standard simplex method implementation. When solving a range of larger Netlib problems and very dense machine learning problems, speed-ups of between 1.6 and 1.8 were achieved when doubling the number of processors. They also presented results which indicated that the performance of their implementation was generally superior to the Minos 5.4 serial sparse revised simplex solver, particularly for the machine learning and denser Netlib problems. Thomadakis and Liu [67] also used steepest-edge in their implementation of the standard simplex method on a MasPar MP-1 and MP-2. Solving a range of large, apparently randomly-generated problems, they achieved a speed-up of up to three orders of magnitude on the  $128 \times 128$  processor MP-2. Further theoretical and practical work on massively parallel implementations of the standard simplex method with steepest-edge has also been reported by Yarmish [71].

## 4.2 Parallel simplex using sparse matrix algebra

The real challenge in developing a parallel simplex implementation of general practical value is to exploit parallelism when using efficient sparse matrix algebra techniques. Only then could the resulting solver be competitive with a good serial implementation when solving general large sparse LP problems using a realistic number of processors.

At the time when the parallelisation of the simplex method was first and most widely considered, practical parallel factorisation and solution methods for sparse unsymmetric linear systems were in their infancy. As a consequence, although work on parallelising the simplex method using dense matrix algebra has generally exploited data parallelism, it was a commonly held view that when using sparse matrix algebra, there was relatively little scope for exploiting data parallelism (with the exception of PRICE). This has led some to conclude that there is no real scope for developing a worthwhile parallel implementation. However, despite the apparently sequential nature of the computational components of the revised simplex method, there is scope for exploiting task parallelism.



### 4.2.1 Sparse standard simplex method

It was observed above that the standard simplex method is generally implemented using dense matrix algebra. However, based on the observation that fill-in in the tableau can be low, Lentini *et al.* [48] developed a parallel implementation of the standard simplex method with the tableau stored as a sparse matrix. In this remarkably novel approach, they incorporated a Markowitz-like criterion into column and row selection with the aim of reducing fill-in. When solving medium sized Netlib problems on four transputers they achieved a speed-up of between 0.5 and 2.7, with a super-linear speed-up of 5.2 on one problem with a relatively large column-row ratio. They compared their results on eight transputers with a sparse revised simplex solver running on a mainframe (without giving any justification of the validity of this comparison) and found that their parallel solver was a little faster on smaller problems but slower on larger problems, particularly when solving four proprietary problems. Although they describe their results as disappointing, this brave and imaginative approach to parallelising the simplex method deserves to be better known.

For the LP problems considered by Lentini *et al.* , the average density of the optimal tableau was found to be about 10%. For large problems, such a density would make the serial sparse standard simplex method prohibitively expensive and only with a prohibitively large number of parallel processes could it conceivably compete with a good serial revised simplex solver. However, Hall and McKinnon [38] have identified that for problems that exhibit hyper-sparsity, the tableau density can be expected to be of order 1% and significantly lower still for some problems. Although this level of sparsity still corresponds to a huge memory overhead for large LP problems, the use of the parallel sparse simplex method to study, rather than solve, hyper-sparse LP problems is the subject of current and projected research by Hall, Hamilton and McKinnon.

### 4.2.2 Revised simplex method with a sparse inverse

When the revised simplex method with an explicit inverse is implemented, dense data structures are generally used to maintain  $B^{-1}$ . However, Shu [64] identified that  $B^{-1}$  is by no means full and considered parallelising the revised simplex method with  $B^{-1}$  maintained using a sparse data structure. At first sight, her results using up to 64 processors on a Touchstone Delta and up to 16 processors on an Intel iSPC/2 hypercube appear good. Using the sparse representation of  $B^{-1}$ , as well as the traditional dense representation, Shu obtained a speed-up of up to 17 on the Touchstone Delta when solving small

to modest Netlib problems. Unfortunately, for all but one of these problems the average density of  $B^{-1}$  is between 33% and 47% so there is little difference in the solution times when using a sparse or dense  $B^{-1}$ . Of those problems for which comparative dense-sparse results are given, only in the case of SHIP08L is the average sparsity (6%) of  $B^{-1}$  worth exploiting: the sparse implementation was twice as fast as the dense implementation and a speed-up of 16 was achieved, compared with a speed-up of 17 for the dense case. Results for significantly larger problems, including some of the CRE and OSA problems from the Kennington test set [17] are also given although, unfortunately, it is not clear how  $B^{-1}$  was represented. Speed-ups of up to 8 were achieved on both machines. However, for these larger problems, analysis of the computational cost of components of Shu's implementation of the revised simplex method with a factored inverse [64] indicated that PRICE, which parallelises easily, constitutes between 65% and 96% of the solution time. Although these proportions would be lower in the explicit inverse implementation, they are much higher than would be the case were Shu to have updated reduced costs and use a row-wise PRICE, exploiting the moderate to high levels of hyper-sparsity that these problems are now known to possess [38]. If the results for the problems where PRICE dominates are excluded, the speed-ups obtained are in low single figures, even when using up to 64 processors.

### 4.2.3 Revised simplex method with a factored inverse

The general inefficiency of alternative means of implementing the simplex method in serial is such that the only real scope for developing a worthwhile parallel simplex solver is to identify how the revised simplex method with a factored inverse may be parallelised. The natural data parallelism of PRICE has been exploited by most authors. Some have then considered the data parallelism in other computational components, whereas others have studied the extent to which task parallelism can be exploited by overlapping operations that are then executed either in serial or by exploiting data parallelism.

### Theoretical pure data parallel approaches

In an early work, Pfefferkorn and Tomlin [61] discussed how parallelism could be exploited in each computational component of the revised simplex method. This purely data-parallel scheme, devised for the ILLIAC IV, was never implemented. After identifying the parallelism of a column PRICE operation, they then considered FTRAN and BTRAN. They identified that if a sequence  $E_s^{-1} \dots E_r^{-1}$  within the factored representation of  $B^{-1}$  has the property that

the pivots are in distinct rows and the eta vectors have no entries in these rows, then the application of these etas is independent in both FTRAN and BTRAN. This may be illustrated, without loss of generality, by assuming that  $r = 1$ ,  $s = K$  and that the pivots are in rows  $1, \dots, K$ , in which case it is readily established that

$$E_K^{-1} \dots E_1^{-1} = \begin{bmatrix} I & \\ E & I \end{bmatrix} \begin{bmatrix} M & \\ & I \end{bmatrix}$$

where  $E_{ij} = -\eta_{K+i}^j$  and  $M = \text{diag}\{\mu^1, \dots, \mu^K\}$ . The corresponding operations  $E_K^{-1} \dots E_1^{-1} \mathbf{r}$  in FTRAN may now be performed as

$$\begin{aligned} \mathbf{r}_K &:= M \mathbf{r}_K \\ \mathbf{r}_{K'} &:= \mathbf{r}_{K'} + E \mathbf{r}_K \end{aligned}$$

where  $\mathbf{r} = \begin{bmatrix} \mathbf{r}_K \\ \mathbf{r}_{K'} \end{bmatrix}$  and, in BTRAN,  $\mathbf{r}_K^T := (\mathbf{r}_K^T + \mathbf{r}_{K'}^T E) M$ . Note that for this technique to be effective it is necessary for most of the (initial) components of  $\mathbf{r}_K$  to be nonzero or, equivalently, that most of the etas will have to be applied in the conventional view of FTRAN and BTRAN. For LP problems exhibiting hyper-sparsity this does not occur so this technique will not be of value. For other problems, the assumption that such sequences exist to a significant extent may be unreasonable. Any techniques developed by Pfeifferkorn and Tomlin for exploiting parallelism during INVERT, which is based on that of Tomlin [68], are either not stated or unclear, possibly referring to the potential for parallelising the search for singletons in the triangularisation phase.

Helgason *et al.* [40] discussed the scope for parallelising FTRAN and BTRAN based on the quadrant interlocking factorisation of Evans and Hatzopoulos [25]. This relatively unknown factorisation scheme corresponds to forming an  $LU$  decomposition of (a row and column permutation of)  $B$  in which, along their diagonals,  $L$  has  $2 \times 2$  blocks and  $U$  has  $2 \times 2$  identity matrices. Thus the entire eta file following INVERT consists of sequences of the type identified by Pfeifferkorn and Tomlin [61], each of length two. It follows that if  $\mathbf{r}$  fills in sufficiently fast during FTRAN and BTRAN, there is the potential for speed-up bounded above by two. Developed with a shared memory machine in mind, this scheme was never implemented. Whilst the potential speed up of two is small, even if it were actually achieved, an extension of this scheme using larger blocks could have greater potential.

McKinnon and Plab [55] considered how data parallelism could be exploited in FTRAN for both dense and sparse right hand sides. They also investigated [54] how the Markowitz criterion could be modified in order to

increase the potential for exploiting data parallelism in subsequent FTRAN operations.

Since the time when the parallelisation of the simplex method was first and most actively considered, there have been huge efforts and advances in practical parallel factorisation and solution methods for sparse unsymmetric linear systems. However, the techniques thus developed have not yet been applied to the revised simplex method, the scope for which is discussed in Section 5.

### **Parallel PRICE for the dual simplex method**

For the (primal) revised simplex method, it is highly questionable to claim the value of an implementation that exploits parallelism in only the PRICE operation on the basis that LP problems with very large column/row ratios exist. As identified above, if such problems are solved with the revised simplex method, then a serial implementation using partial pricing is likely to be very much more efficient. However, the same is not true when solving LP problems using the *dual* revised simplex method. This algorithm, which is very much more efficient for many single LP problems and is valuable when solving LP subproblems in branch-and-bound for integer programming, has computational requirements in each iteration that are very similar to those of the primal algorithm. Where it differs fundamentally is that there is no scope for restricting PRICE to just a subset of the nonbasic variables. Bixby and Martin [14] investigated the scope for parallelism in the dual simplex method and chose to parallelise only those operations whose cost is related to the number of columns in the problem, that is PRICE, the dual ratio test and update of the dual variables. Much the most expensive of these is, naturally, PRICE which, in the context of the dual algorithm, is used to refer to the matrix-vector product  $\mathbf{z}^T N$ . Each process formed a part of  $\mathbf{z}^T N$ , with the corresponding columns of  $N$  held row-wise so that sparsity in  $\mathbf{z}$  can be exploited. The work of Bixby and Martin is also notable since the underlying dual simplex code being parallelised was CPLEX. Thus any speed-up was necessarily relative to a state-of-the-art serial solver. They implemented the dual simplex method on several architectures but, for the full Netlib set, the overhead of exploiting parallelism exceeded any gain in performance. This is unsurprising since few problems in the Netlib set have significantly more columns than rows. Focusing on test problems with very large column/row ratios, using two Sun workstations connected by ethernet they achieved very little speed-up. Performance using up to 4 processors on an IBM SP2 was better, with speed-up ranging from 1 to 3. Switching to experiments using 2 processors on shared memory machines, little speed-up


was obtained with a Sun S20-502 but better results were obtained with a Silicon Graphics multiprocessor. In later work reported in this paper [14], the speed-up achieved using 4 processors of a shared memory SGI Power challenge was better still, being near linear for most problems with more than 64 columns per row. Strangely this behaviour tailed off for the problems with most extreme numbers of columns per row.

### Task and data parallel revised simplex

The first attempt to exploit task parallelism in the revised simplex method was reported by Ho and Sundarraaj [43]. In addition to the natural data parallelism of (column-wise) PRICE, Ho and Sundarraaj identified that INVERT can be overlapped with simplex iterations. Once formed, the representation of  $B_0^{-1}$  is generally not up-to-date with respect to basis changes that occurred during INVERT. However, when using the product form update, a representation of  $B^{-1}$  is readily obtained by incorporating the eta vectors formed since INVERT was started. By dedicating one processor solely INVERT, Ho and Sundarraaj observe that the frequency of INVERT should be increased. This can be expected to reduce the cost of FTRAN and BTRAN since the representation of  $B^{-1}$  will involve fewer eta vectors. Whilst they claim that having fewer eta vectors will have advantages in terms of numerical stability, they do not discuss the potential for numerical instability identified earlier by Hall and McKinnon [33]. As can be expected for only a partially parallelised algorithm, the practical performance of Ho and Sundarraaj's implementation, on an Intel iPSC/2 and Sequent Balance 8000, is limited in accordance with Amdahl's law. On a set of small Netlib and proprietary problems, they report an average saving of 33% over the serial solution time.

To date, there have only been three implementations of the revised simplex method that have attempted to exploit parallelism anything like fully: one by Wunderling [69, 70] and two by Hall and McKinnon [34, 35]. Wunderling's variant of the revised simplex method was based on multiple pricing with steepest edge weights. The theoretical task parallelism is conveniently illustrated by the Gantt chart in Figure 4.

Wunderling did not perform explicit multiple pricing since all steepest edge weights were updated after each basis change. However, since the remaining tableau columns for attractive candidates were updated, the variant was equivalent to multiple pricing. Since INVERT was not overlapped with other calculations, the implementation was not quite fully parallel. If this serial operation is excluded then the implementation is fully parallel for two processors so long as there is not a significant load imbalance between the two BTRAN operations. However, a good simplex solver will exploit the fact



C: CHUZC R: CHUZR P1: Update PRICE P2: Steepest edge PRICE

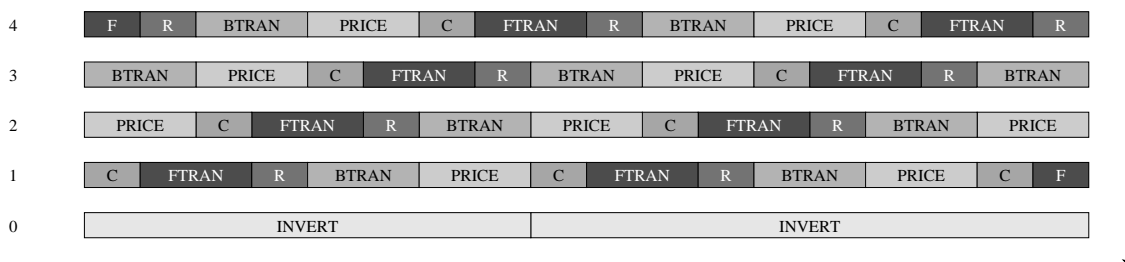
Figure 4: Wunderling’s parallel revised simplex with steepest edge

that the right hand side of the standard ‘update’ BTRAN,  $\pi_p^T = e_p^T B^{-1}$ , contains only one nonzero (in a pivotal row), whereas it may not be worth exploiting any sparsity in the right hand side of the additional steepest edge BTRAN,  $w = \hat{a}_q^T B^{-1}$ . Thus the parallelism of the BTRAN operations may be illusory.

It is unfortunate that the only published account of Wunderling’s most interesting work on parallel (and serial) simplex is his PhD Thesis [69] which is written in German. At a workshop presentation [70], he described his parallel simplex implementation and gave limited results. Wunderling’s parallel scheme was implemented on a 2-processor Sun 20 and a Cray T3D using modules of Soplex, his public-domain simplex solver which (still) outperforms other such solvers in tests performed by Mittelman [57]. On general LP problems any performance improvement that was achieved was minimal and significant speed-up was only achieved for problems with large column/row ratios.

The first of Hall and McKinnon’s two parallel revised simplex schemes was ASYNPLEX [35]. This scheme corresponds to a variant of the revised simplex method in which reduced costs are computed directly (as  $\hat{c}_N^T = c_N^T - \pi_B^T c_B^T B^{-1} N$ ) rather than being updated and Dantzig pricing is used. Ideally, one processor is devoted to INVERT and the remaining processors perform simplex iterations. The theoretical operations of these processors are illustrated by the Gantt chart in Figure 5. If a single path through the feasible region is to be followed then, clearly, two processors cannot perform basis changes simultaneously. An additional processor was introduced to act as ‘basis change manager’ to ensure that a processor is only allowed to start CHUZR if its basis is up-to-date and no other processor has been allowed to start CHUZR at the current basis. Although each processor has reduced costs, more up-to-date reduced costs may be known on other processors. Thus, once a processor has computed the full set of reduced costs, a small set of good candidates are communicated to a further

processor. This acts as a ‘column selection manager’ and responds to requests made prior to FTRAN by the iteration processors, by communicating the best (known) candidate to enter the basis. Note that, due to the fact that basis changes were determined during the course of calculating reduced costs, column selection was generally performed using out-of-date reduced costs. These reduced costs became further out-of-date as a consequence of basis changes occurring during FTRAN.



C: CHUZYC R: CHUZR

Figure 5: ASYNPLEX

ASYNPLEX was implemented on a Cray T3D using modules of Hall’s highly efficient revised simplex solver whose performance, on the four modest but representative Netlib test problems used for experiments, was comparable with that of OSL [44]. Using between eight and twelve iteration processors, the speed with which simplex iterations were performed was increased by a factor of about five in all cases. However, the increase in the number of iterations required to solve the problem (resulting from the use of out-of-date reduced costs) led to the speed-up in solution time ranging from 2.4 to 4.8. Although the speed-up obtained was felt to be satisfactory, it was limited by various factors. A further consequence of limited candidate persistence was the time wasted on candidates which, upon completion of FTRAN, were found to be unattractive. The overheads of communicating the pivotal column from the iteration processor upon which it was computed, to all other iteration processors, together with the duplicated work of bringing the eta file up-to-date on all processors, also reduced efficiency. Numerical instability, due to the need to re-use eta vectors from the previous representation of  $B^{-1}$ , was observed for many test problems.

Hall and McKinnon’s second parallel scheme was PARSMI [34]. This was developed in an attempt to address the deficiencies of ASYNPLEX [35]. In order to make realistic comparisons with good serial solvers, PARSMI updates the reduced costs and uses Devex pricing. To reduce the communication overheads, each processor is associated with one of the major computational

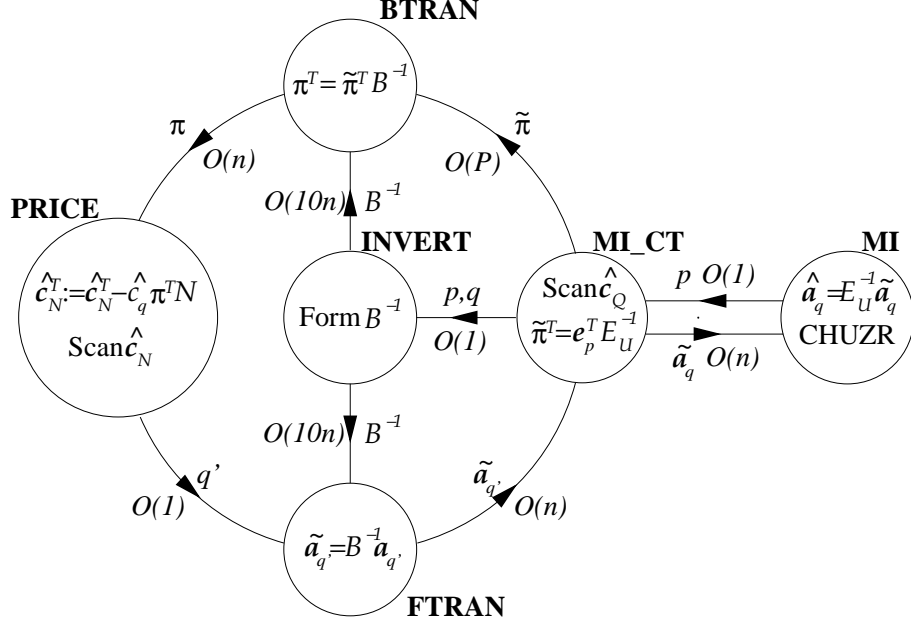


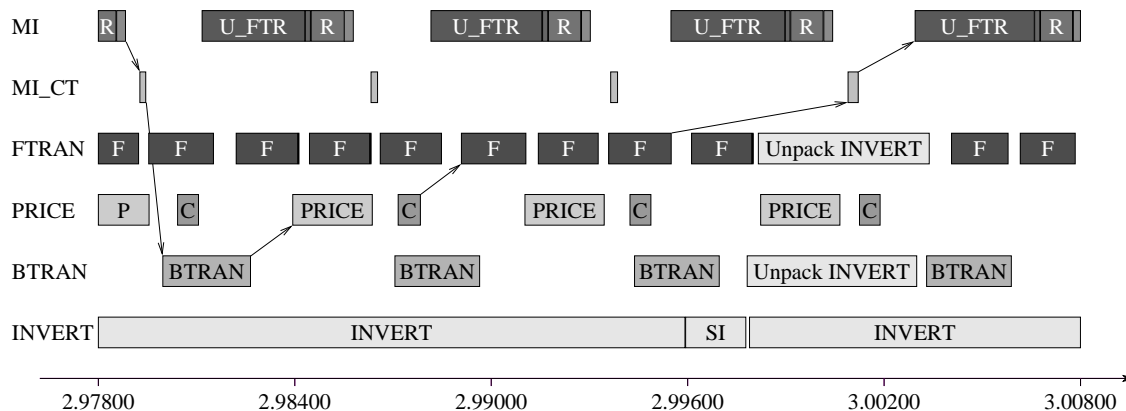
Figure 6: Operations and inter-processor communications for PARSMI

components of the revised simplex method, with a variant of multiple pricing being used to increase the scope for parallelism. The operations performed on processors of different types and the inter-processor communications are illustrated in Figure 6, where MI represents the processors performing minor iterations and MLCT is the single processor coordinating the MI processors.

The FTRAN processors receive sets of attractive candidates to enter the basis from the PRICE processors and apply the INVERT etas. The resulting partially completed tableau columns are distributed over the MI processors where the update etas are applied. The MLCT processor determines the best candidate for which the up-to-date tableau column is known and coordinates CHUZR, which is distributed over the MI processes. Once the pivotal row has been determined, the MLCT processor applies the update etas to start BTRAN. This is done very efficiently using the technique described by Hall and McKinnon [38]. The resulting partially formed vector  $\pi_p$  is then communicated to a BTRAN processor which applies the INVERT etas before distributing the completed vector  $\pi_p$  across the PRICE processors. After a distributed PRICE operation, the PRICE processors update the reduced costs and Devex weights before determining sets of candidates and communicating them to the FTRAN processors. One or more INVERT processors factorise the basis matrix serially and communicate the result to the FTRAN and BTRAN processors. Although presented sequentially, operations are overlapped as



much as possible. For a six-processor implementation, the operation of the scheme is illustrated in Figure 7. Note that this Gantt chart was created from actual data when applying the T3D implementation of PARSMI to the LP test problem 25FV47.



C: CHUZR C: CHUZR F: I-FTRAN U\_FTR: U-FTRAN R: CHUZR SI: Send INVERT

Figure 7: Real processor activity and illustrative communication for PARSMI

As with ASYNPLEX, the implementation of PARSMI was a highly complex programming task, magnified by fact that communication times are not deterministic and the order of arrival of messages determined the operation of the program. Programming difficulties, together with numerical instability (for the same reason as with ASYNPLEX) meant that the implementation was never reliable. In the very limited results that were obtained on modest Netlib problems, using eight processors the speed-up in iteration speed of between 2.2 and 2.4 was eroded by the consequences of limited candidate persistence to yield a speed-up in solution time of between 1.7 and 1.9.

#### 4.2.4 The simplex method on supercomputers

From the results quoted by Beasley [10] in a review of experience with the simplex method on Cray supercomputers, it may be concluded that these machines were used principally to study algorithmic performance when solving large LP problems, rather than to investigate the scope for exploiting their vector arithmetic facilities. When the latter was not used, the performance decreased by only a few tens of percent.

The most notable attempt to exploit a vector facility is reported by Forrest and Tomlin [28]. They vectorised (column-wise) PRICE by identifying those sets of columns of the constraint matrix for which the numbers of

nonzeros are equal and at least 20, and storing the values and row indices in dense rectangular arrays. Using a set of moderate test problems, some of which were drawn from the Netlib set, they achieved speed-ups that are generally modest (less than 2), although for one problem the speed-up of 3.8 compares well with the figure of 4.5 which is the maximum attainable on the particular architecture. They did not consider the consequences of  $\pi$  being sparse, in which case their vectorised PRICE would involve many operations with zero. For a sparse  $\pi$ , a row-oriented PRICE avoids operations with zero but, since the numbers of nonzeros in  $N$  change from one iteration to another, it is hard to imagine their vectorisation technique being of value. Forrest and Tomlin go on to consider the extent to which a vector processor can be exploited in FTRAN and BTRAN, as well as UPDATE. In their serial revised simplex implementation, the latter has a significant computational overhead since it modifies the factors obtained by INVERT but, typically, yields a significantly smaller eta file than would be obtained with the product form update. Estimated speed-ups of 1.5 and 2.0 for FTRAN, BTRAN and UPDATE are reported for two illustrative problems. As for other computational components, Forrest and Tomlin observe that when packing the pivotal column prior to CHUZR in ‘easy’ problems (where the pivotal column contains a large proportion of zeros) much time is spent identifying nonzeros and identify that this operation vectorises well. In modern revised simplex solvers, exploiting such hyper-sparsity during FTRAN leads to list of indices of (potential) nonzeros being known so the full packing operation is unnecessary. The overall effect of exploiting a vector processor is a general speed-up of up to 4, with a speed-up of 12 for one problem, relative to a (then) state-of-the-art revised simplex solver.

#### 4.2.5 Following parallel paths

An alternative approach to exploiting parallel computation when using the simplex method is to allow each processor to follow different paths through the feasible region and, periodically, determine which processor has made the best progress and then repeat the process from the corresponding basis. The simplest possible scheme of this kind would correspond to a parallelisation of the *maximum improvement* rule for column selection. This is a criterion in which CHUZR is performed for all attractive candidates to enter the basis and the basis change performed is that which yields the greatest reduction in the objective. Unless it is performed within the context of the standard simplex tableau, as a variant for serial computation it is prohibitively expensive. However, it parallelises naturally. A more practical approach that has been considered is to distribute the best candidates, one to each processor,

and allow the simplex method to perform a fixed number of basis changes before determining the process that has made the best progress. Whilst this is simple to implement and has minimal communication overhead, it cannot be expected to yield worthwhile speed-up. One argument why is that, if the simplex method starts from a basis of logical variables and the optimal basis consists of structural variables, then the number of simplex iterations that must be performed is bounded below by  $m$ . Since efficient serial simplex solvers will, typically, determine the optimal basis in a small multiple of  $m$  iterations, this is an immediate bound on the speed-up that can be achieved. For some LP problems, typically those that are highly degenerate, the number of iterations required even by good solvers is a significant multiple of  $m$ . However, even for such problems, it is doubtful that the number of iterations required would be reduced sufficiently by any parallel path-following implementation.

Another approach to following parallel paths has been investigated by Maros and Mitra [53]. Rather than using the same variant of the revised simplex method to follow paths, Maros and Mitra's strategy is to start from different edges at a common vertex and use eight different pricing strategies, including Dantzig, Devex, multiple pricing and several variants of partial pricing. They refer to this as a parallel *mixed* strategy. Although Devex pricing is commonly used as a default strategy, it is well known that for certain LP problems it is far from being optimal. This is particularly so for problems with large column/row ratios and Maros and Mitra give evidence of this. Further, it is also widely acknowledged that, during the course of solving a given LP problem, the best pricing strategy is likely to change. Early in the solution process, when the basis matrix contains a high proportion of logical variables, Devex may be significantly less effective than partial or multiple pricing. Many efficient serial simplex solvers have a run-time pricing strategy that first decides whether the column/row ratio is such that partial pricing it is expected to be more efficient and, if not, starts with some cheap pricing strategy before switching to a full pricing strategy such as Devex.

Maros and Mitra implemented their mixed strategy on a twelve processor Parsytec CC-12, using an extension of their efficient FortMP simplex solver and tested it on fifteen of the largest Netlib problems. Note that the default pricing strategy for FortMP is dynamic cyclic partial multiple pricing. The mixed strategy is generally more efficient than the best single serial strategy (by an average factor of 1.19), and always more efficient than both the FortMP default strategy (by an average factor of 1.69) and Devex (by an average factor of 1.56). Whilst the the mixed strategy was more than three times faster than both the default strategy and Devex (for a different single LP problem in both cases), the average speed-up is modest. It would be

interesting to compare the performance of the mixed strategy with a serial run-time strategy such as that described above.

### 4.3 Parallelising the network simplex method

When the simplex method is applied to a minimum cost network flow problem, exploiting its structure yields the network simplex method. Although it exploits sparsity and structure very efficiently, the data structures and algorithmic requirements for implementing the network simplex method are much simpler than those when implementing the revised simplex method for general LP problems. The parallelisation of the network simplex method has been considered by several authors, including Chang *et al.* [18], Peters [60] and Barr and Hickman [9]. The latter two are of interest since they identify the value of overlapping computational components, subject to sufficient co-ordination to ensure that the simplex method follows a single path through the feasible region. As such they predate the task-parallel approaches to the revised simplex method of Hall and McKinnon [34, 35] and Wunderling [69, 70].

### 4.4 Parallelising simplex-like methods

The parallelisation of other simplex-like methods for solving LP problems have also been studied. These techniques, whilst generally uncompetitive with a good implementation of the revised simplex method, offer more immediate scope for parallelisation. They are also frequently only of relevance to LP problems which possess a particular, but common, structure. This is either inherent, as a consequence of the underlying model, or identified by analysing the constraint matrix.

In the technique of Dantzig-Wolfe decomposition [20] for appropriate block-angular problems, the requirement to solve a (possibly large) number of independent smaller LP subproblems is a natural source of parallelism. This was identified by Ho *et al.* [42] and load-balancing issues were addressed by Gnanendran and Ho [31]. Boduroglu [15] parallelised the little-known technique of Kaul [45] for block-angular problems and obtained some impressive results, even beating a world class revised simplex solver on several standard test problems. The solution of problems with block-angular structure using parallel bundle methods is described by Medhi [56].

For block-angular LP problems whose structure is the transpose of that required for Dantzig-Wolfe decomposition, Benders decomposition [11] may be used. The requirement to solve (nested) sets of independent LP problems has been exploited in a number of parallel implementations, for example [23,

41, 58]. Rosen and Mayer [63] describe a method for block-angular problems with Dantzig-Wolfe structure in which, after forming the dual problem, a procedure similar to Benders decomposition yields a set of independent LP subproblems to be solved in parallel.

Work on identifying block angular structure in general LP problems was also reported by Pinar and Aykanat [62] and these techniques are still being developed by Aykanat *et al.* [7]. Ferris and Horn [26] discuss how block-angular structure may be identified in general LP problems prior to applying parallel bundle methods, for which they achieved an efficiency of approximately 90% when using 32 processors of a Thinking machines CM-5.

Klabjan *et al.* [46] describe a primal-dual simplex algorithm in which several primal subproblems are solved in parallel. Their algorithm is only efficient for problems with very high column/row ratios and, for such problems, they achieve a speed-up of up to 4 on a cluster of 16 PCs.

## 4.5 Summary

Attempts to exploit parallelism in the simplex method have been associated with many of the leading figures in the development of efficient serial techniques for the revised simplex method. That the success relative to good serial implementations has been limited is a measure of the difficulty of the task.

None of the attempts to exploit parallelism in the simplex method has, for general large sparse LP problems, offered significantly improved performance over a good serial implementation of the revised simplex method. Parallelising the standard or revised simplex method with dense matrix algebra is uncompetitive unless a prohibitive number of processors is used. With sparse matrix algebra, parallelising just the PRICE operation, even for LP problems with a high column/row ratio, is uncompetitive with the revised simplex method if partial pricing is used. For problems with a high column/row ratio, if the dual simplex method is competitive with the revised simplex method with partial pricing, then parallelising PRICE will offer a worthwhile performance improvement. The success of attempts to exploit task parallelism has been limited by lack of candidate persistence, numerical instability and the communication overheads of distributed memory multi-processors.

## 5 Directions for the future

Since the period when most of the work on parallelising the simplex method took place, there have been advances in the availability of shared memory multiprocessors, as well as parallel factorisation and solution techniques for general unsymmetric sparse linear systems. As a consequence of both of these developments, work in the area of parallelising the revised simplex method with sparse matrix algebra merits renewed consideration.

### 5.1 Exploiting task parallelism

Of the three implementations of the revised simplex method that have attempted to exploit parallelism anything like fully, the more successful were ASYNPLEX and PARSMI. Of the three limitations of these schemes, numerical instability and communication overhead are currently being addressed by Hall. The former is eliminated by overlapping INVERT with multiple BTRAN, multiple PRICE, CHUZC and multiple FTRAN. In an attempt to reduce the latter, this scheme is being implemented on a shared memory Sun Fire E15k. The limitations due to lack of candidate persistence are expected to be reduced but cannot be eliminated.

### 5.2 Exploiting data parallelism

To date, there has been no published attempt to develop a fully data-parallel implementation of the revised simplex method using sparse matrix algebra. A principal reason for this is that, at the time of the greatest activity in the parallelisation of the simplex method, there had been relatively little progress in parallel methods for factorising and solving general sparse linear systems of equations.

Since this time, there has been considerable progress in the development and implementation of parallel techniques for the factorisation of general sparse matrices, for example SuperLU from Berkeley [22, 49] and MUMPS from Lyon-Toulouse [4, 5]. Many features distinguish the factorisation and linear system solution requirements in the revised simplex method. The fact that the matrix being factored is often highly reducible can be exploited in many ways. It has been shown by Amestoy *et al.* [3, 6] that general preprocessing and reordering techniques can combine numerical accuracy and reducibility detection in the context of general sparse solvers. Further, since the right hand sides of FTRAN and BTRAN are sparse, the reducibility of the matrix may also be used to exploit hyper-sparsity and therefore accelerate the parallel solution phase of general sparse solvers. With the fruits of these

advances as yet un-tapped in the context of the revised simplex method, the time is right for more work in this area.

The author would like to thank Patrick Amestoy for his insight into the developments in techniques for parallel factorisation and solution techniques for general unsymmetric sparse linear systems, and the potential for their application to the revised simplex method.

## References

- [1] A. Agrawal, G. E. Blelloch, R. L. Krawitz, and C. A. Phillips. Four vector-matrix primitives. In *ACM Symposium on Parallel Algorithms and Architectures*, pages 292–302, 1989.
- [2] G. M. Amdahl. Validity of the single-processor approach to achieving large scale computing capabilities. In *AFIPS Conference Proceedings*, volume 30, pages 483–485. AFIPS Press, Reston, Va., 1967.
- [3] P. Amestoy, S. Pralet, and S. Li. Unsymmetric orderings using a constrained Markowitz scheme, February 27-28 2004. SIAM conference on Parallel Processing for Scientific Computing, Workshop of Combinatorial Scientific Computing, San Francisco (California).
- [4] P. R. Amestoy, I. S. Duff, J. Koster, and J.-Y. L’Excellent. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM Journal on Matrix Analysis and Applications*, 23(1):15–41, 2001.
- [5] P. R. Amestoy, I. S. Duff, and J.-Y. L’Excellent. Multifrontal parallel distributed symmetric and unsymmetric solvers. *Comput. Methods Appl. Mech. Eng.*, 184:501–520, 2000.
- [6] P. R. Amestoy, X. S. Li, and S. Pralet. Constrained Markowitz with local symmetrization. Technical Report RT/APO/04/05, ENSEEIHT-IRIT, Toulouse, France, 2004. Also appeared as Lawrence Berkeley Lab report LBNL-56861 and CERFACS report TR/PA/04/137, submitted to SIAM Journal on Matrix Analysis and Applications.
- [7] C. Aykanat, A. Pinar, and Ü. V. Çatalyürek. Permuting sparse rectangular matrices into block-diagonal form. To appear in SIAM Journal on Scientific Computing.

- [8] D. A. Babaev and S. S. Mardanov. A parallel algorithm for solving linear programming problems. *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki*, 31(1):86–95, 1991.
- [9] R. S. Barr and B. L. Hickman. Parallel simplex for large pure network problems: Computational testing and sources of speedup. *Operations Research*, 42(1):65–80, 1994.
- [10] J. E. Beasley. Linear programming on Cray supercomputers. *Journal of the Operational Research Society*, 41(2):133–139, 1990.
- [11] J. F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4:238–252, 1962.
- [12] R. E. Bixby. Solving real-world linear programs: A decade and more of progress. *Operations Research*, 50(1):3–15, 2002.
- [13] R. E. Bixby, M. Fenelon, Z. Gu, E. Rothberg, and R. Wunderling. MIP: Theory and practice closing the gap. In M. J. D. Powell and S. Scholtes, editors, *System Modelling and Optimization: Methods, Theory and Applications*, pages 19–49. Kluwer, The Netherlands, 2000.
- [14] R. E. Bixby and A. Martin. Parallelizing the dual simplex method. *INFORMS Journal on Computing*, 12:45–56, 2000.
- [15] İ. İ. Boduroğlu. *Scalable Massively Parallel Simplex Algorithms for Block-Structured Linear Programs*. PhD thesis, GSAS, Columbia University, New York, NY, 1997.
- [16] T. B. Boffey and R. Hay. Implementing parallel simplex algorithms. In *CONPAR 88*, pages 169–176, Cambridge, UK, 1989. Cambridge University Press.
- [17] W. J. Carolan, J. E. Hill, J. L. Kennington, S. Niemi, and S. J. Wichmann. An empirical evaluation of the KORBX algorithms for military airlift applications. *Operations Research*, 38(2):240–248, 1990.
- [18] M. D. Chang, M. Engquist, R. Finkel, and R. R. Meyer. A parallel algorithm for generalized networks. *Annals of Operations Research*, 14:125–145, 1988.
- [19] Z. Cvetanovic, E. G. Freedman, and C. Nofsinger. Efficient decomposition and performance of parallel PDE, FFT, Monte-Carlo simulations, simplex, and sparse solvers. *Journal of Supercomputing*, 5:19–38, 1991.



- [20] G. B. Dantzig. The decomposition principle for linear programs. *Operations Research*, 8:101–111, 1960.
- [21] G. B. Dantzig and W. Orchard-Hays. The product form for the inverse in the simplex method. *Math. Comp.*, 8:64–67, 1954.
- [22] J. W. Demmel, S. C. Eisenstat, J. R. Gilbert, X. S. Li, and J. W. H. Liu. A supernodal approach to sparse partial pivoting. *SIAM Journal on Matrix Analysis and Applications*, 20(3):720–755, 1999.
- [23] M. A. H. Dempster and R. T. Thompson. Parallelization and aggregation of nested Benders decomposition. *Annals of Operations Research*, 81:163–187, 1998.
- [24] J. Eckstein, İ. İ. Boduroğlu, L. Polymenakos, and D. Goldfarb. Data-parallel implementations of dense simplex methods on the Connection Machine CM-2. *ORSA Journal on Computing*, 7(4):402–416, 1995.
- [25] D. J. Evans and M. Hatzopoulos. A parallel linear system solver. *International Journal of Computer Mathematics*, 7:227–238, 1979.
- [26] M. C. Ferris and J. D. Horn. Partitioning mathematical programs for parallel solution. *Mathematical Programming*, 80:35–61, 1998.
- [27] R. A. Finkel. Large-grain parallelism—three case studies. In L. H. Jamieson, D. Gannon, and R. J. Douglas, editors, *The Characteristics of Parallel Algorithms*, pages 21–63. MIT Press, Cambridge, MA, 1987.
- [28] J. J. H. Forrest and J. A. Tomlin. Vector processing in the simplex and interior methods for linear programming. *Annals of Operations Research*, 22:71–100, 1990.
- [29] D. M. Gay. Electronic mail distribution of linear programming test problems. *Mathematical Programming Society COAL Newsletter*, 13:10–12, 1985.
- [30] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright. A practical anti-cycling procedure for linearly constrained optimization. *Mathematical Programming*, 45:437–474, 1989.
- [31] S. K. Gnanendran and J. K. Ho. Load balancing in the parallel optimization of block-angular linear programs. *Mathematical Programming*, 62:41–67, 1993.

- [32] D. Goldfarb and J. K. Reid. A practical steepest-edge simplex algorithm. *Mathematical Programming*, 12:361–371, 1977.
- [33] J. A. J. Hall and K. I. M. McKinnon. Update procedures for the parallel revised simplex method. Technical Report MSR 92-13, Department of Mathematics and Statistics, University of Edinburgh, 1992.
- [34] J. A. J. Hall and K. I. M. McKinnon. PARSMI, a parallel revised simplex algorithm incorporating minor iterations and Devex pricing. In J. Waśniewski, J. Dongarra, K. Madsen, and D. Olesen, editors, *Applied Parallel Computing*, volume 1184 of *Lecture Notes in Computer Science*, pages 67–76. Springer, 1996.
- [35] J. A. J. Hall and K. I. M. McKinnon. ASYNPLEX, an asynchronous parallel revised simplex method algorithm. *Annals of Operations Research*, 81:27–49, 1998.
- [36] J. A. J. Hall and K. I. M. McKinnon. Exploiting hyper-sparsity in the revised simplex method. Technical Report MS99-014, Department of Mathematics and Statistics, University of Edinburgh, 1999.
- [37] J. A. J. Hall and K. I. M. McKinnon. The simplest examples where the simplex method cycles and conditions where EXPAND fails to prevent cycling. *Mathematical Programming*, 100(1):133–150, 2004.
- [38] J. A. J. Hall and K. I. M. McKinnon. Hyper-sparsity in the revised simplex method and how to exploit it. To appear in *Computational Optimization and Applications*, 2005.
- [39] P. M. J. Harris. Pivot selection methods of the Devex LP code. *Mathematical Programming*, 5:1–28, 1973.
- [40] R. V. Helgason, L. J. Kennington, and H. A. Zaki. A parallelisation of the simplex method. *Annals of Operations Research*, 14:17–40, 1988.
- [41] R. S. Hiller and J. Eckstein. Stochastic dedication: designing fixed income portfolios using massively parallel Benders decomposition. *Management Science*, 39(11):1422–1438, 1993.
- [42] J. K. Ho, T. C. Lee, and R. P. Sundarraaj. Decomposition of linear programs using parallel computation. *Mathematical Programming*, 42:391–405, 1988.

- [43] J. K. Ho and R. P. Sundarraj. On the efficacy of distributed simplex algorithms for linear programming. *Computational Optimization and Applications*, 3(4):349–363, 1994.
- [44] IBM. *Optimization Subroutine Library, guide and reference, release 2*, 1993.
- [45] R. N. Kaul. An extension of generalized upper bounding techniques for linear programming. Technical Report ORC 65-27, O. R. Center U. C. Berkley, San Fransisco, CA, 1965.
- [46] D. Klabjan, E. L. Johnson, and G. L. Nemhauser. A parallel primal-dual simplex algorithm. *Operations Research Letters*, 27:47–55, 2000.
- [47] V. Kumar, A. Grama, A. Gupta, and G. Karypis. *Introduction to Parallel Computing: Design and Analysis of Algorithms*. Addison-Wesley, 2 edition, 2003.
- [48] M. Lentini, A. Reinoza, A. Teruel, and A. Guillen. SIMPAR: a parallel sparse simplex. *Computational and Applied Mathematics*, 14(1):49–58, 1995.
- [49] X. S. Li and J. W. Demmel. SuperLU\_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems. *ACM Transactions on Mathematical Software*, 29(2), 2003.
- [50] J. Luo and G. L. Reijns. Linear programming on transputers. In J. van Leeuwen, editor, *Algorithms, Software, Architecture*, volume A-12 of *IFIP Transactions A (Computer Science and Technology)*, pages 525–534. Elsevier, 1992.
- [51] J. Luo, G. L. Reijns, F. Bruggeman, and G. R. Lindfield. A survey of parallel algorithms for linear programming. In E. F. Deprettere and A.-J. van der Veen, editors, *Algorithms and Parallel VLSI Architectures*, volume B, pages 485–490. Elsevier, 1991.
- [52] H. Markowitz. The elimination form of the inverse and its application to linear programming. *Management Science*, 3:255–296, 1957.
- [53] I. Maros and G. Mitra. Investigating the sparse simplex algorithm on a distributed memory multiprocessor. *Parallel Computing*, 26:151–170, 2000.

- [54] K. I. M. McKinnon and F. Plab. A modified Markowitz criterion to increase parallelism in inverse factors of sparse matrices. Technical report, Department of Mathematics and Statistics, University of Edinburgh, 1997.
- [55] K. I. M. McKinnon and F. Plab. An upper bound on parallelism in the forward transformation within the revised simplex method. Technical report, Department of Mathematics and Statistics, University of Edinburgh, 1997.
- [56] D. Medhi. Parallel bundle-based decomposition algorithm for large-scale structured mathematical programming problems. *Annals of Operations Research*, 22:101–127, 1990.
- [57] H. D. Mittelmann. Benchmarks for optimization software. <http://www.plato.la.asu.edu/bench.html>, April 2002.
- [58] S. N. Nielsen and S. A. Zenios. Scalable Benders decomposition for stochastic linear programming. *Parallel Computing*, 23:1069–1088, 1997.
- [59] W. Orchard-Hays. *Advanced Linear programming computing techniques*. McGraw-Hill, New York, 1968.
- [60] J. Peters. The network simplex method on a multiprocessor. *Networks*, 20:845–859, 1990.
- [61] C. E. Pfefferkorn and J. A. Tomlin. Design of a linear programming system for the ILLIAC IV. Technical Report SOL 76-8, Systems Optimization Laboratory, Stanford University, 1976.
- [62] A. Pinar and C. Aykanat. An effective model to decompose linear programs for parallel solution. In J. Waśniewski, J. Dongarra, K. Madsen, and D. Olesen, editors, *Applied Parallel Computing*, volume 1184 of *Lecture Notes in Computer Science*, pages 592–601. Springer, 1996.
- [63] J. B. Rosen and R. S. Maier. Parallel solution of large-scale, block-angular linear programs. *Annals of Operations Research*, 22:23–41, 1990.
- [64] W. Shu. Parallel implementation of a sparse simplex algorithm on MIMD distributed memory computers. *Journal of Parallel and Distributed Computing*, 31(1):25–40, November 1995.
- [65] C. B. Stunkel. Linear optimization via message-based parallel processing. In *International Conference on Parallel Processing*, volume III, pages 264–271, August 1988.

- [66] U. H. Suhl and L. M. Suhl. Computing sparse LU factorizations for large-scale linear programming bases. *ORSA Journal on Computing*, 2(4):325–335, 1990.
- [67] M. E. Thomadakis and J.-C. Liu. An efficient steepest-edge simplex algorithm for SIMD computers. In *International Conference on Supercomputing*, pages 286–293, 1996.
- [68] J. A. Tomlin. Pivoting for size and sparsity in linear programming inversion routines. *J. Inst. Maths. Applics*, 10:289–295, 1972.
- [69] R. Wunderling. Paralleler und objektorientierter simplex. Technical Report TR-96-09, Konrad-Zuse-Zentrum für Informationstechnik Berlin, 1996.
- [70] R. Wunderling. Parallelizing the simplex algorithm. ILAY Workshop on Linear Algebra in Optimization, Albi, April 1996.
- [71] G. Yarmish. *A Distributed Implementation of the Simplex Method*. PhD thesis, Polytechnic University, Brooklyn, NY, March 2001.